Random curves on surfaces induced from the Laplacian determinant

Adrien Kassel * Richard Kenyon †

November 30, 2012

Abstract

We define natural probability measures on cycle-rooted spanning forests (CRSFs) on graphs embedded on a surface with a Riemannian metric. These measures arise from the Laplacian determinant and depend on the choice of a unitary connection on the tangent bundle to the surface.

We show that, for a sequence of graphs $(\mathcal{G}_n)$ conformally approximating the surface, the measures on CRSFs of $\mathcal{G}_n$ converge and give a limiting probability measure on finite multicurves (finite collections of pairwise disjoint simple closed curves) on the surface, independent of the approximating sequence.

Wilson's algorithm for generating spanning trees on a graph generalizes to a cycle-popping algorithm for generating CRSFs for a general family of weights on the cycles. We use this to sample the above measures. The sampling algorithm, which relates these measures to the loop-erased random walk, is also used to prove tightness of the sequence of measures, a key step in the proof of their convergence.

We set the framework for the study of these probability measures and their scaling limits and state some of their properties.

Contents

*Département de Mathématiques et Applications, École Normale Supérieure, 45 rue d'Ulm, 75230 Paris Cedex 05, France. Partially supported by Fondation Sciences Mathématiques de Paris.

†Mathematics Department, Brown University, 151 Thayer St., Providence, RI 02912. Research supported by the NSF.

1 Introduction

A *cycle-rooted spanning forest* (CRSF) on a graph $\mathcal{G}$ is a subgraph each of whose components contains a unique cycle, or equivalently, contains as many vertices as edges, see Figure 1. A *cycle-rooted spanning tree* (CRST) is a connected CRSF. Natural probability measures on CRSFs arising from the determinant of the graph Laplacian were introduced in [8]: the probability of a CRSF is proportional to the product of its edge weights times a product over its cycles of a certain function of the cycle, depending on the monodromy of a discrete $\mathbb{C}^*$- or $\mathrm{SL}_2(\mathbb{C})$-connection.

The interest of these measures is that they give to a cycle a weight which is a function of its shape (homotopy type or enclosed curvature in the topological or geometrical settings, respectively, as described below).

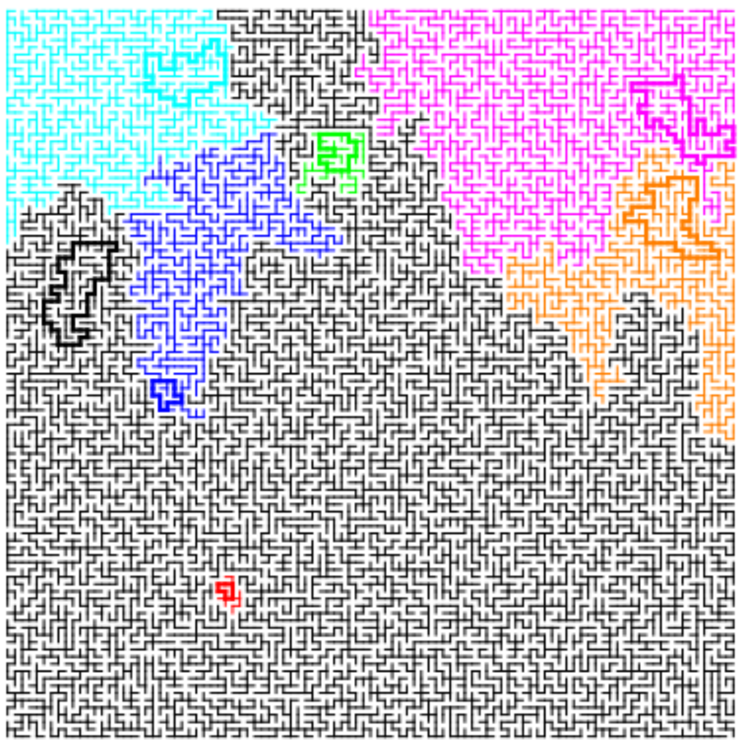

Figure 1: A CRSF on the 100×100 square grid; each connected component is in a different color and the cycles are in bold.

We give a "cycle-popping" algorithm (Theorem 1) for rapid exact sampling from these and other measures, generalizing the well-known cycle-popping algorithm of

Wilson [11] for generating uniform spanning trees. One simply runs Wilson's algorithm, and when a cycle is created, flip a coin (with bias depending on the cycle weight) to decide whether to keep it or not.

We use this sampling algorithm to sample approximations of scaling limits of CRSFs on surfaces in two different settings: a topological and a geometrical one, as follows. By *surface* we will mean here an oriented smooth surface with a Riemannian metric. By *approximation* of these, we will mean a sequence of graphs $(\mathcal{G}_n)$ geodesically embedded on a surface Σ and *conformally approximating* Σ in a sense defined below (essentially, the simple random walk on $\mathcal{G}_n$ converges to Brownian motion on Σ). We endow the space of multiloops with a natural topology and show the weak convergence of probability measures on discrete multiloops to probability measures on continuous multiloops as follows.

In the topological setting, we consider only the conformal class of the metric. An *incompressible* CRSF is a CRSF with no contractible cycles. Let $(\mathcal{G}_n)$ be an approximating sequence and μ^n_{inc} be the uniform measure on incompressible CRSFs of $\mathcal{G}_n$. We show (Theorem 15) that the cycle process $\mathbb{P}^n_{inc}$ of the μ^n_{inc}-random CRSF on $\mathcal{G}_n$ converges to a random loop process $\mathbb{P}_{inc}$ on Σ, independent of the sequence $\mathcal{G}_n$. This only depends on the conformal class of Σ, in the following sense. Let $z_1, \dots, z_k$ be distinct points of Σ, and $\delta > 0$ be small; then for any isotopy class of sets of pairwise disjoint simple closed loops $\{\gamma_1, \dots, \gamma_m\}$ of $\Sigma \setminus \{B_\delta(z_1) \cup \dots \cup B_\delta(z_k)\}$, the probability that a random incompressible CRSF on $\mathcal{G}_n$ has m cycles, and these are isotopic to the γ_i (which we call a *cylindrical event*), has a probability converging as $n \to \infty$ to a limit independent of the approximating sequence $\mathcal{G}_n$. This refines the result of [8] which shows that the distribution of the homotopy classes of the cycles in $\pi_1(\Sigma)$ has a conformally invariant limit.

In the geometrical setting, we take into account the metric on Σ, and in particular its curvature. Let $(\mathcal{G}_n)$ be a sequence of finite graphs conformally approximating Σ. Associated to this data is a discrete connection on a complex line bundle over $\mathcal{G}_n$ arising from the Levi-Civita connection on the tangent bundle on Σ. It is defined up to gauge equivalence by the property that the monodromy around a closed loop is $e^{i\theta}$ where θ is the enclosed curvature. From this connection Φ^n_{LC} we construct a natural probability measure μ^n_{LC} on CRSFs on $\mathcal{G}_n$: each CRSF has a probability proportional to the product of its edge weights times the product over its cycles of $2 - 2\cos\theta$, where θ is the curvature enclosed. The corresponding loop process $\mathbb{P}^n_{LC}$ is shown (Theorem 17) to converge to a probability measure $\mathbb{P}_{LC}$ on multicurves on Σ.

When the surface Σ is contractible, we define another measure $\mu^n_{LC^0}$ which is in some sense more natural. This measure is a limit of the CRSF connection measures with curvature $e^{i\varepsilon\theta}$ when $\varepsilon \to 0$, a limit which was introduced in [8]. This yields

a measure on CRSTs (CRSFs with one component) with weight proportional to θ^2 where θ is the enclosed curvature. We show (Theorem 17) that the loop measure $\mathbb{P}^n_{LC^0}$ converges to a measure $\mathbb{P}_{LC^0}$ on simple closed curves on Σ.

In all three cases, our generalization of Wilson's algorithm is used to show tightness of the sequence of measures (Section 4.3). Convergence follows from the convergence of the measures on a determining class, the above-mentioned cylindrical events (Lemma 13) in the flat case and, in addition, convergence of polygonal approximations in the curved case.

We give samples from the measures μ_{LC} and μ_{LC^0} for the round sphere (Figures 2, 3), a saddle surface (Figure 4) and a compact disk in the Poincaré plane (Figures 5), and for the measure μ_{inc} on a flat torus (Figure 6) and planar domains (Figures 7 and 9). For μ_{LC} these are conditional samples, conditioned on having only loops with area (curvature) bounded by $\pi/2$; the above sampling algorithm does not work without this condition (see, however, [6] where it is shown how to sample from any determinantal process with Hermitian kernel, of which μ_{LC} is one).

The paper is organized as follows. In Section 2 we introduce the sampling algorithm and prove its correctness. In Section 3 we introduce the probability measures on CRSFs on graphs on surfaces and show how they are exactly sampled by the algorithm. In Section 4 we show that the probability measures on loops that these induce converge to probability measures on the space of multiloops of the surface. Section 5 enumerates some of the properties of the measures on loops considered in the paper. The paper concludes with a list of open questions in Section 6.

2 A general sampling algorithm

An *oriented* CRSF is a CRSF in which each cycle has a chosen orientation. A measure on CRSFs induces a measure on oriented CRSFs by giving each cycle an independent $1/2-1/2$-chosen orientation, and a measure on oriented CRSFs induces one on CRSFs by forgetting the orientation.

Let $\mathcal{G} = (V, E)$ be a finite graph with vertex and edge sets V and E, respectively, and $c : E \to \mathbb{R}_{>0}$ a positive function on the edges which we call the *conductance*. Let α be a function which assigns to each oriented simple loop γ in $\mathcal{G}$ a positive weight $\alpha(\gamma) \in [0, 1]$. These functions c, α define a probability measure $\mu = \mu_{c,\alpha}$ on oriented CRSFs, giving an oriented CRSF Γ a probability proportional to $\prod_{e\in\Gamma} c(e) \prod_{\text{cycles } \gamma\subset\Gamma} \alpha(\gamma)$. We describe an algorithm to sample an oriented CRSF according to the measure μ.

We note that this sampling algorithm requires $\alpha \in [0, 1]$; it will not work without modification for larger α. In the special case where $\alpha = 1$ and $c = 1$, the algorithm

samples according to the uniform measure on oriented CRSFs.

Let us describe a cycle-popping procedure, named $P[w, \Gamma]$, which takes as arguments w a vertex and Γ an oriented subgraph of $\mathcal{G}$ not containing w, and outputs another oriented subgraph of $\mathcal{G}$ containing Γ and w. The procedure is the following: start at vertex w and perform a simple random walk (with each step proportional to the conductances) until it first reaches a vertex v which either belongs to Γ or is the first self-intersection of its path;

- If v is in Γ, then replace Γ by the union of Γ and the oriented path just traced by the random walk.
- If v is the first self-intersection, let γ denote the oriented cycle thus obtained, and sample a $\{0,1\}$-Bernoulli random variable with success probability $\alpha(\gamma)$;
  - If the outcome is 1, then replace Γ by the union of Γ and the oriented path just traced by the random walk.
  - If the outcome is 0, erase the cycle that was just closed and continue to perform the random walk from v until it reaches Γ or self-intersects, in which case repeat the above instructions.

The algorithm, called $\mathcal{A}$, is then the following: start with Γ empty and w an arbitrary vertex in V, perform $P[w, \Gamma]$ until Γ contains all vertices of $\mathcal{G}$, and output Γ.

Theorem 1. *If $\alpha(\gamma) > 0$ for some γ then the algorithm $\mathcal{A}$ terminates and its output is an oriented CRSF, sampled according to the measure μ.*

Remark 2. *Note that if in the above definition of $\mathcal{A}$ one starts with Γ equal to S, a distinguished set of vertices in $\mathcal{G}$, then the algorithm samples an oriented "essential" CRSF with Dirichlet boundary conditions on S, see [8] or section 5.1 below for a definition.*

Proof. As in Wilson's proof, we construct an equivalent description, denoted $\mathcal{A}'$, of the algorithm $\mathcal{A}$.

Let us define $\mathcal{A}'$ in the following way. Consider, over each vertex $v \in V$, an infinite sequence $X^{(v)}$ of i.i.d random variables $X^{(v)} = (X_1^{(v)}, X_2^{(v)}, \dots)$, each distributed as a random neighbor of v according to the conductance measure, that is, for each $i \geq 1$ and each neighbor w of v, we have

$$\mathbb{P}\left(X_i^{(v)} = w\right) = \frac{c(vw)}{\sum_{v' \sim v} c(vv')} .$$

We represent $X^{(v)}$ as an infinite stack of cards, with only $X_1^{(v)}$ being visible at the top of the stack.

We draw an edge from each vertex v to the neighbor shown on the top of the stack $X^{(v)}$; the oriented graph thus seen is an oriented CRSF (with possible loops of length 2). This is our initial CRSF. We now describe a step by step random popping algorithm of the cycles. Note that at each step, the graph that we see remains an oriented CRSF. Here is the algorithm: For each cycle γ encountered in the current CRSF, pop it with probability proportional to $1-\alpha(\gamma)$; when a cycle is popped off, the top card on the stacks for each of its vertices is discarded. When a cycle is "kept", its cards are fixed and can no longer be removed. The algorithm stops once the cycles that remain have been all "kept" in a Bernoulli trial. It is easy to see that the order in which the cycles are popped is not relevant.

Note that this algorithm terminates (as long as there is at least one cycle with positive probability) because eventually, with probability 1, all the cycles present at one step will have been previously kept.

Since the card of the stacks are distributed as the steps of a conductance-biased random walk, we see that algorithm $\mathcal{A}'$ has the same output in distribution as algorithm $\mathcal{A}$. In order to compute the output distribution of $\mathcal{A}$ we will therefore use algorithm $\mathcal{A}'$.

Let us compute the probability that a given oriented CRSF Γ is obtained as an output of algorithm $\mathcal{A}'$. Let $\gamma_1, \ldots, \gamma_k$ be the cycles of Γ. The CRSF Γ is obtained as an output if and only if there exists a finite sequence of oriented cycles $C_1, \cdots, C_m$ such that these cycles are popped, and after removing them, the cards that appear correspond to Γ, and there are k successful trials for Bernoullis with success probability $\alpha(\gamma_i)$.

By independence of the cards in the stacks, the last CRSF considered is independent of the cycles that were popped. Therefore, for any oriented CRSF Γ, we have

$$\begin{aligned}
\mathbb{P}(\Gamma) &= \sum_{\mathcal{C}=\{C_1,\ldots,C_m\}} \mathbb{P}\left(\Gamma \mid \text{pop } \mathcal{C}\right)\mathbb{P}\left(\text{pop } \mathcal{C}\right) \\
&= \sum_{\mathcal{C}} \mathbb{P}\left(\text{pop } \mathcal{C}, \text{ and } \Gamma \text{ occurs underneath and is kept}\right) \\
&= \sum_{\mathcal{C}} \mathbb{P}\left(\text{pop } \mathcal{C}\right) \prod_{e\in\Gamma} \mathbb{P}(e) \prod_{i=1}^{k} \alpha(\gamma_i) \\
&= \left(\sum_{\mathcal{C}} \mathbb{P}\left(\text{pop } \mathcal{C}\right)\right) \prod_{e\in\Gamma} \mathbb{P}(e) \prod_{i=1}^{k} \alpha(\gamma_i)
\end{aligned}$$

which we see is proportional to the weight of Γ. $\square$

To sample a non-oriented CRSF according to a measure which assigns a CRSF Γ a weight proportional to $\prod_{e\in\Gamma} c(e)\prod_{\gamma\subset\Gamma}\alpha(\gamma)$, where the product is over non-oriented cycles γ, and α is a function invariant under orientation, it suffices to have $\alpha\in[0,2]$, perform algorithm $\mathcal{A}$ for the measure $\mu_{c,\alpha/2}$, and forget the orientation in the resulting oriented CRSF. In particular, we obtain the uniform measure on non-oriented CRSFs with the choice $c=\alpha=1$.

There is a variant of the previous algorithm to sample an oriented CRSF according to measure $\mu_{c,\alpha}$ conditional on having a single loop: multiply all the loop weights by a small constant ε. Then perform $\mathcal{A}$; if ε is small there will typically be a single loop (if not, start over).

Let N be the total number of vertices of the graph. The running time of the algorithm is bounded by the time to obtain the first loop (which is bounded by $O(N^2)$ if $\alpha\geq O(1/N^2)$) plus the running time of Wilson's algorithm, that is $O(N(\log N)^2)$ (Wilson's algorithm has a running time bounded by the cover time [11] which is linear up to a logarithmic correction [1]). The running time is at least linear. Extreme cases correspond to extreme values of α: for $\alpha=1$ (uniform measure on CRSFs), the running time is linear; for α close to zero (like in the conditional measure described in the previous paragraph), the running time is large, at least $O(1/\sup\alpha)$.

3 Natural probability measures on CRSFs

The most natural probability measure on CRSFs on a finite graph is the uniform measure. There are however other natural probability measures that can be constructed from connections on line bundles and that are meaningful for graphs embedded in surfaces.

3.1 Connections

Let $\mathcal{G}=(V,E)$ be a finite graph. A *complex line bundle on* $\mathcal{G}$ is a copy $\mathbb{C}_v$ of $\mathbb{C}$ associated to each vertex $v\in V$. The total space of the bundle is the direct sum $W=\bigoplus_{v\in V}\mathbb{C}_v$. A *unitary connection* Φ on W is the data consisting of, for each oriented edge $e=vv'$, a unitary complex linear map $\varphi_{vv'}:\mathbb{C}_v\to\mathbb{C}_{v'}$ such that $\varphi_{v'v}=\varphi_{vv'}^{-1}$ (equivalently, we assign to each oriented edge e a unit-modulus complex number φ_e such that $\varphi_{-e}=\overline{\varphi}_e$). We say that two connections Φ,Φ' are *gauge equivalent* if there exist unit-modulus complex numbers ψ_v such that $\psi_{v'}\varphi_{vv'}=\varphi'_{vv'}\psi_v$, that is, Φ' is obtained from Φ by changing the basis of each space $\mathbb{C}_v$ by a rotation. Let

$c : E \to \mathbb{R}_{>0}$ be a conductance function. We let Δ_Φ be the associated Laplacian acting on $f \in W$ defined, for each vertex v, by

$$\Delta_\Phi(f)(v) = \sum_{v' \sim v} c(vv') \left(f(v) - \varphi_{v'v} f(v')\right) ,$$

where the sum is over all neighbors v' of v.

When $\mathcal{G}$ is geodesically embedded in a surface Σ (that is, embedded in such a way that edges are geodesic segments), there is a natural connection $\Phi = \Phi_\nabla$ on $\mathcal{G}$ arising from a connection ∇ on the tangent bundle $T\Sigma$: we define for each vertex v the line $\mathbb{C}_v$ to be the tangent plane to Σ, with its natural complex structure coming from the metric. The ∇-parallel transport along edge e of $\mathcal{G}$ defines the parallel transport ϕ_e.

3.2 Laplacian determinant and measures

Theorem 3 ([5, 8]). *For a graph with connection Φ on a complex line bundle we have*

$$\det(\Delta_\Phi) = \sum_{\text{CRSFs}} \prod_{edges} c(e) \prod_{cycles} (2 - 2\cos\theta) ,$$

where $e^{i\theta}$ is the monodromy of the connection around the cycle, for any choice of its orientation.

Associated to Φ is a probability measure μ_Φ on CRSFs, where the probability of a CRSF is proportional to $\prod_{\text{edges}} c(e) \prod_{\text{cycles}} (2 - 2\cos\theta)$.

Another measure, introduced in [8], is obtained by passing to the limit $t \to 0$ in a connection with weights $\varphi_e = e^{it\theta_e}$. This yields a measure on CRSTs (one-component CRSFs) which assigns a weight θ^2 to each CRST with cycle γ, where $e^{i\theta}$ is the monodromy around γ (for any choice of its orientation) of the initial connection at $t = 1$.

3.3 Flat connections

Suppose that $\mathcal{G}$ is geodesically embedded on a surface Σ with flat connection ∇ on its tangent bundle. Let $\Phi = \Phi_\nabla$ be the associated connection on $\mathcal{G}$. Since ∇ is flat, it has no curvature around contractible cycles. So the associated measure μ_Φ gives zero weight to contractible cycles and thus is supported on incompressible CRSFs.

Let μ_{inc} be the *uniform* measure on incompressible CRSFs on $\mathcal{G}$. The measure μ_Φ has density $\prod_{\gamma \subset \Gamma} (2 - 2\cos\theta_\gamma)$ with respect to μ_{inc}.

Although μ_{inc} cannot itself be written as a connection-measure μ_Φ for some flat connection ∇, we can use the μ_Φ to study μ_{inc}, see Lemma 13 below.

3.4 Graphs embedded on a curved surface

3.4.1 The Levi-Civita measure

Suppose that $\mathcal{G}$ is geodesically embedded on a curved surface Σ. Take ∇ to be the Levi-Civita connection associated to a metric g on Σ. Define μ_{LC} to be the associated probability measure. It gives a CRSF a probability proportional to $\prod_{\gamma\subset\Gamma}(2-2\cos\theta_\gamma)$ where θ is the curvature enclosed by γ.

3.4.2 The CRST measure

As discussed above, there is another measure μ_{LC^0} we can associate to this situation, but only when Σ is contractible. It is supported on the CRSTs of $\mathcal{G}$. Let $\Phi = \{e^{i\theta_e}\}_{e\in E}$ be the parallel transports on $\mathcal{G}$ defined from ∇, and let $\Phi_t = \{e^{it\theta_e}\}_{e\in E}$; these are well defined by contractibility of Σ. Let μ_{LC^0} be the limit as $t \to 0$ of the measures μ_{Φ_t}. In μ_{LC^0}, each CRST Γ has a weight proportional to θ_γ^2, the square of the curvature enclosed by the unique cycle γ of Γ.

Let $Z_{LC^0} = \sum_{\text{CRSTs}} \theta_\gamma^2$ be the partition function of μ_{LC^0}. By Theorem 3, we have

$$Z_{LC^0} = \lim_{t\to 0} t^{-2} \det \Delta_{\Phi_t}\,.$$

The partition function can also be computed as follows. Let $\kappa_{\mathcal{G}}$ be the number of spanning trees of $\mathcal{G}$ and T the transfer impedance function (see section 4.1.1 for the definition of transfer impedance).

Theorem 4 ([7] Lemma 8)**.** *Let E^+ be an orientation of the edges. We have*

$$Z_{LC^0} = \kappa_{\mathcal{G}} \sum_{e\in E^+} \left(\theta_e - \sum_{e'\in E^+} \theta_e \theta_{e'} T(e,e') \right)\,.$$

3.4.3 CRST and LERW

The measure μ_{LC^0} has a density with respect to a weighted "bridge" measure on loop-erased random walk excursions on the graph. Let $\mathbb{P}_{LC^0}$ be the measure on simple loops in $\mathcal{G}$ obtained from μ_{LC^0} by ignoring the branches of the CRST. Define $\mathbb{P}_{\text{LERW}}$ to be the law of a loop-erased random walk (LERW) from x to y where $e = xy$ is a uniform random oriented edge in $\mathcal{G}$. For a simple loop γ, define $\mathbb{P}_{\text{LERW}}(\gamma)$ to be the probability that (1) the uniform random edge xy is an edge of γ, and (2) the LERW from x to y is the remainder of γ.

Lemma 5. *For any cycle γ in $\mathcal{G}$, we have*

$$\mathbb{P}_{LC^0}(\gamma) = \frac{|E|\kappa_{\mathcal{G}}}{Z_{LC^0}} \frac{\theta_\gamma^2}{|\gamma|} \mathbb{P}_{\mathrm{LERW}}(\gamma)\,,$$

where $|E|$ is the number of edges of $\mathcal{G}$.

Proof. We have

$$\begin{aligned}
\mathbb{P}_{LC^0}(\gamma)\, Z_{LC^0} &= \sum_{\mathrm{CRST}\supset\gamma} \theta_\gamma^2 \\
&= \frac{\theta_\gamma^2}{|\gamma|} \sum_{e\in\gamma} |\{\text{spanning trees which contain } \gamma\setminus e\}| \\
&= \frac{\theta_\gamma^2}{|\gamma|} |E|\kappa_{\mathcal{G}} \mathbb{P}_{\mathrm{LERW}}(\gamma)\,.
\end{aligned}$$

□

3.5 Exact sampling

The measures μ_{inc}, μ_{LC^0}, and μ_{LC} can be sampled using our generalization of Wilson's algorithm as follows.

The measure μ_{inc} is sampled by using a function α which assigns a loop weight 0 if it is contractible and 1 otherwise.

For μ_{LC^0}, we set $\alpha(\gamma) = \varepsilon\theta_\gamma^2$ for small ε. For small enough ε there will typically be only one loop (if there is more than one, start over).

We can sample μ_{LC} only in the case where the absolute value of the curvature θ enclosed by any curve doesn't exceed $\pi/2$. Indeed, in that case, we will always have $2 - 2\cos\theta \in [0, 2]$ which is necessary to sample.

See Figures 2, 3, 4, 5, 6 which are obtained by using fine conformal approximations to the underlying surfaces (only the cycles of the CRSFs are drawn). Figures 3 and 5 are samples conditional on enclosing curvature less than $\pi/2$.

In order to sample the unconditional measures, one can use a general algorithm for determinantal processes with Hermitian kernels [6] which is slower. Figures 7 and 9 show samples of μ_{inc} on multiply-connected planar domains.

4 Scaling limits for graphs on surfaces

The measures $\mu_{inc}, \mu_{LC}, \mu_{LC^0}$ induce measures on sets of cycles on $\mathcal{G}$, by forgetting the rest of the CRSF. We use notation $\mathbb{P}_{inc}, \mathbb{P}_{LC}, \mathbb{P}_{LC^0}$ respectively for these cycle

Figure 2: A μ_{LC^0}-random CRSF on the sphere with its round metric

Figure 3: A μ_{LC}-random CRSF on the sphere with two components

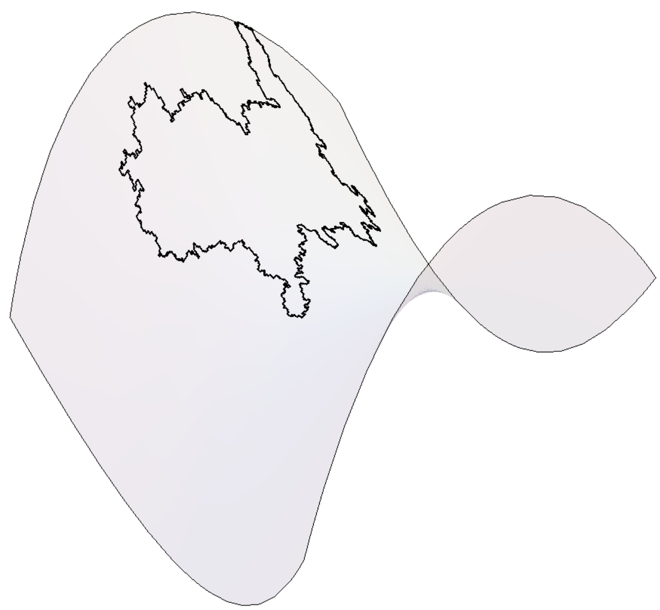

Figure 4: A μ_{LC}-random CRSF on on the saddle surface $z = x^2 - y^2$

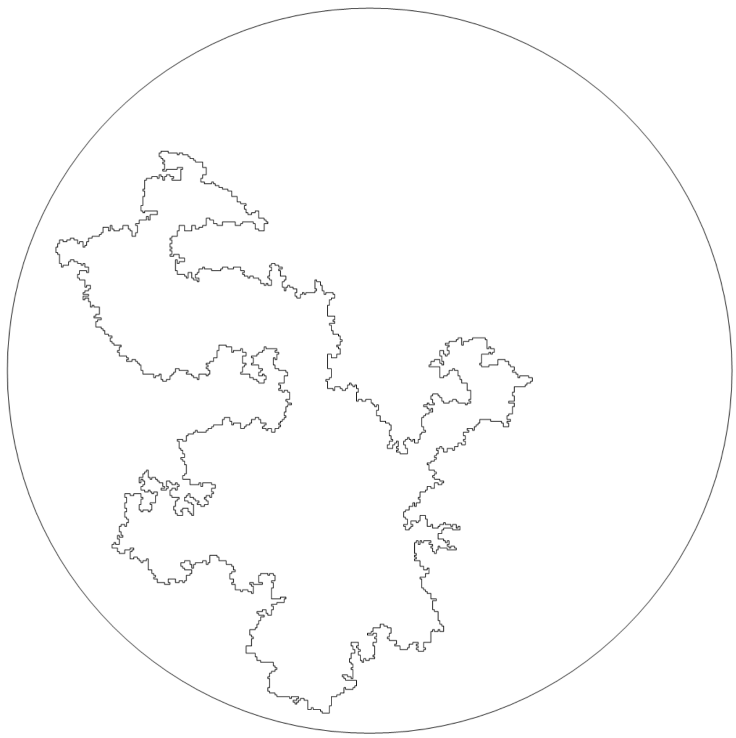

Figure 5: A μ_{LC}-random CRSF in a hyperbolic ball of radius 2 centered at $(0, 1)$ in the hyperbolic plane

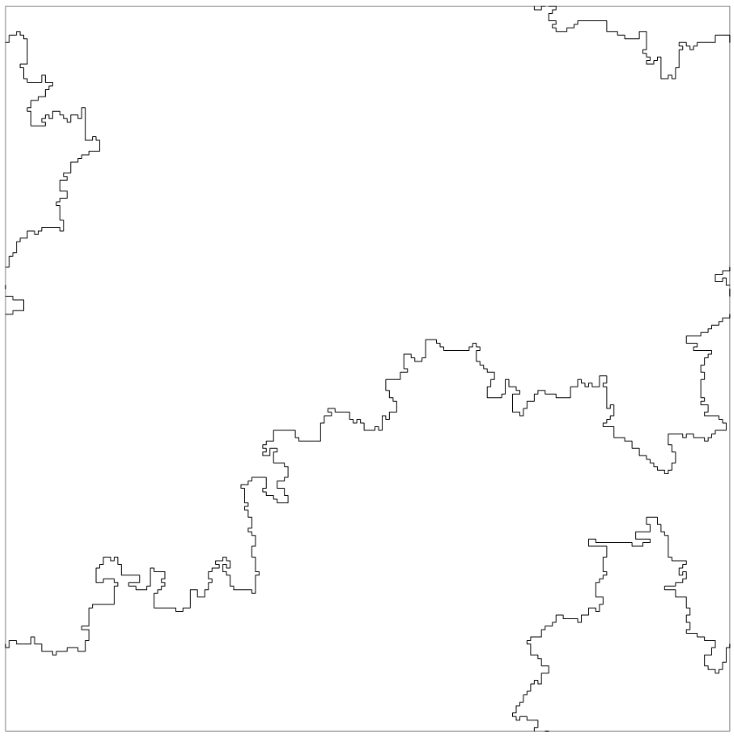

Figure 6: A μ_{inc}-random CRSF on the flat torus (which is obtained from the unit square by identifying opposite sides)

measures.

4.1 Conformal Approximation

Let $(\mathcal{G}_n)$ be a sequence of (edge-weighted) graphs geodesically embedded in Σ with mesh size (longest edge length) going to zero.

There are a number of equivalent definitions of the notion of conformal approximation of Σ by the sequence $(\mathcal{G}_n)$. Perhaps the easiest is to say that (conductance-weighted) random walk on $\mathcal{G}_n$ converges to the Brownian motion on Σ, up to reparameterization. Another, more computationally useful, approach uses the transfer impedance.

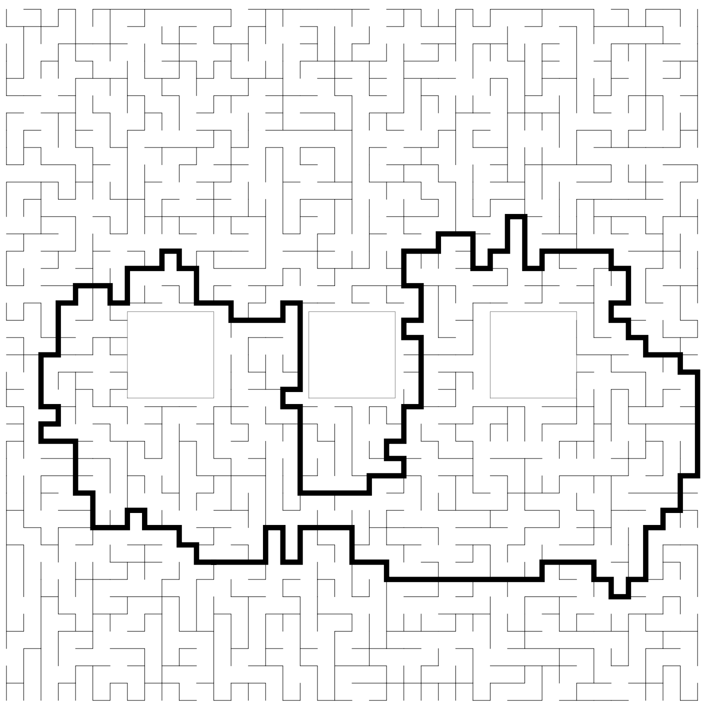

Figure 7: A μ_{inc}-random CRSF on a punctured disk conditioned on having a particular homotopy type

4.1.1 Transfer impedance

For any n let $T^n(e, e')$ be the transfer impedance of two oriented edges $e = xy$ and $e' = x'y'$ in $\mathcal{G}_n$. This is defined as the (algebraic) current through e' when one unit of current enters at e_- and leaves at e_+ (the endpoints of e). In terms of the Green function G^n one has

$$T^n(e, e') = G^n(e_+, e'_+) - G^n(e_+, e'_-) - G^n(e_-, e'_+) + G^n(e_-, e'_-).$$

The function T^n is a discrete one-form on $\mathcal{G}_n$ (function on oriented edges which changes sign under change of orientation). We say that *conformal approximation holds* if when e_1, e_2 are not within $o(1)$ of each other,

$$T^n(e_1, e_2) = \left(\frac{\partial}{\partial e_1}\frac{\partial}{\partial e_2} g_D(z_1, z_2)\right)\ell(e_1)\ell(e_2) + o(\ell(e_1)\ell(e_2))\,, \tag{1}$$

where g_D is the continuous Green function, $\ell(e_i)$ is the edge length, and the notation $\frac{\partial}{\partial e_i} g_D$ represents the directional derivative in the direction of the edge e_i applied to z_i.

If we sum (1) for e_1 varying along a path $\gamma_{a,b}$ from vertex a to b, and e_2 on $\gamma_{c,d}$ from c to d, we find that

$$\sum_{e_1 \in \gamma_{a,b}} \sum_{e_2 \in \gamma_{c,d}} G^n(e_1^-, e_2^-) - G^n(e_1^-, e_2^+) - G^n(e_1^+, e_2^-) + G^n(e_1^+, e_2^+) =$$
$$g_D(a, c) - g_D(b, c) - g_D(a, d) + g_D(b, d) + o(1). \tag{2}$$

This quantity represents the change in potential from c to d when one unit of current enters at a and exits at b. As a function $F(c)$ of c, this is the unique harmonic function with appropriate logarithmic singularities at a and b, and which is zero at d. In particular the level curves of F are equipotentials and, taking equipotentials near a and b (which are close to circles), one can thus compute the resistance between a small circle around a and a small circle around b. In this way, using convergence of the transfer impedance, one can show that the main notions of potential theory including the Poisson kernel, Cauchy kernel, holomorphic functions, *etc.* all behave well under conformal approximation.

As an example of a conformally approximating sequence, we can pull back a fine square grid in $\mathbb{C}$ under a conformal map from a domain in $\mathbb{C}$ to Σ. More generally the (almost-)equally-spaced real and imaginary curves of a quadratic differential $\phi(z)dz^2$ on Σ naturally give a graph structure on Σ which is (almost-)conformally equivalent to a square grid graph except for conical singularities (with angles multiples of 2π) at the zeros of ϕ.

4.2 The probability space of simple multiloops

For a graph $\mathcal{G}_n$, let $\mathbb{P}^n$ be one of the measures on loops discussed above. We see the probability measures $\mathbb{P}^n$ as living on the same probability space Ω that we now describe.

A *simple multiloop* in Σ is a finite union of simple loops (ignoring parameterization) with disjoint images, that is, an injective map from the union of k copies of the unit circle $\mathbb{T}$ to Σ, for some $k \geq 1$ (and modulo reparameterizations). The space of single simple loops is a metric space: the distance is defined by

$$d(f,g) = \inf_{\alpha} \sup_{t \in S^1} \|f(t) - g(\alpha(t))\|$$

where the infimum is over all reparameterizations α. In other words two loops are close if they can be parameterized so that their images are close for all parameter values t. One can extend this distance to the case of multiloops by taking the infimum over all permutations of loops, of the pairwise distances (and defining the distance between Ω_k and $\Omega_{k'}$ for $k \neq k'$ to be infinite). With this distance Ω is a topological metric space. It is a disjoint union $\Omega = \cup_{k=1}^{\infty} \Omega_k$ where Ω_k consists of multiloops with k components.

This space is not complete: it is easy to construct Cauchy sequences that shrink to a point or to non self-avoiding loops. However it is separable: take all finite multiloops lying on fine lattice approximations of the surface. This is a countable family of loops which is dense in Ω.

The set of cycles of a CRSF on $\mathcal{G}_n$ defines an element of Ω.

A *finite lamination* on a surface is an isotopy class of a simple multiloop. For any points $x_1, \ldots, x_m \in \Sigma$, and small $\delta > 0$, let B_i be a ball around x_i for some radius less than δ, and consider the finite laminations in $\Sigma \setminus \{B_1 \cup \cdots \cup B_m\}$. Any simple multiloop γ which avoids the balls B_i defines a lamination $[\gamma]$. For any of these laminations L we consider the event

$$E_{B_1,\ldots,B_m;L} = \{\gamma \in \Omega \quad | \quad [\gamma] = L\} \ .$$

We call these sets *cylindrical events* and consider the σ-field $\mathcal{B}$ on Ω generated by these events.

Lemma 6. *$\mathcal{B}$ contains the Borel sets in Ω.*

Proof. We just prove this for one loop, that is, for Ω_1. The proof is easily extended to the general case.

Let c be a smooth simple closed curve in Σ and for some small $\varepsilon > 0$ let $U_\varepsilon(c)$ be its ε-neighborhood. Consider the event E that the random curve γ maps into $U_\varepsilon(c)$,

winding once around the annulus with, say, the positive orientation. These types of events generate the Borel sets.

Let $x_1, x_2, \dots$ be a sequence of points dense in the boundary of $U_\varepsilon(c)$. Let B_i be a ball around x_i of radius δ/i. Let $E_{\delta;n}$ be the cylinder event that the curve γ in $\Sigma \setminus \{B_1 \cup \cdots \cup B_m\}$ separates the points $\{x_1, \dots, x_n\}$ on the two boundary components of $U_\varepsilon(c)$, that is, is consistent with γ winding once around $U_\varepsilon(c)$. The event E is contained in the intersection over n of the $E_{\delta;n}$. In fact we have $E = \cap_n E_{0;n} = \cup_{\delta\to 0} \cap_n E_{\delta;n}$: any continuous simple loop which separates the points lies in the interior of the annulus and winds once around. This can be seen as follows. First of all, any curve in $\cap_n E_{0;n}$ is contained in $U_\varepsilon(c)$: otherwise, we could find some x_j lying on the wrong side of the curve since the family $(x_n)_{n\geq 1}$ is dense in the boundary which would yield a contradiction. Second, the curve cannot be contractible since this would contradict the fact that it separates the points. Hence it winds once around the hole of the annulus. Its orientation is necessarily positive. □

4.3 Asymptotic size of loops and tightness of the measures

4.3.1 Macroscopic loops

For $\mathbb{P}_{inc}$, the loops necessarily are macroscopic since they are noncontractible. In the curved case for the measures $\mathbb{P}_{LC}$ and $\mathbb{P}_{LC^0}$, we show that there are, with positive probability, macroscopic loops.

Theorem 7. *With positive probability $\mathbb{P}_{LC}$ and $\mathbb{P}_{LC^0}$ contain a macroscopic loop, that is, for sufficiently small $\varepsilon > 0$ the probability that there is a loop with diameter $\geq \varepsilon$ does not tend to zero with n.*

Proof. We show this for $\mathbb{P}_{LC^0}$. Absolute continuity between it and $\mathbb{P}_{LC}$ implies that it will hold for $\mathbb{P}_{LC}$ as well.

Consider first the $n \times n$ grid H_n scaled by $1/n$ to the unit square $[0,1]^2$. Let $z_1 \neq z_2 \in (0,1)^2$ and let f_1, f_2 be faces of the grid close to z_1, z_2 respectively. In [7], it is shown that the probability $\mathbb{P}(f_1, f_2)$ that f_1, f_2 are enclosed in the cycle of a uniform CRST of H_n is $const/n^2$. Let E be the event that f_1, f_2 are enclosed in the cycle. On this event, the area of the enclosing cycle is with high probability $\geq \varepsilon n^2$ for some $\varepsilon > 0$ (see below, where it is shown that the loops are locally SLE_2). Then

$$\mathbb{P}_{LC^0}(E) = \frac{\sum_E \mathrm{Area}^2}{\sum_{CRSTs} \mathrm{Area}^2} \geq \frac{\varepsilon^2 n^4 \sum_E 1}{Cn^2 \sum_{CRSTs} 1} = \varepsilon^2 C' n^2 \mathbb{P}(f_1, f_2) \qquad (3)$$

for constants C, C'. Here the denominator of the central inequality follows from the result of [7] that the second moment of the area of a uniform CRST is of order n^2

times a constant. The right-hand side of (3) is bounded by below by a positive constant independently of n.

A similar argument holds for any graph $\mathcal{G}_n$ conformally approximating a surface Σ, and near a point where the curvature of the metric is nonzero, since the arguments of [7] which we used (in both the numerator and denominator above) rely only on the transfer impedance. $\square$

We are not able to show that all loops are macroscopic with probability 1 (although we believe it is true). This appears to rely on a more delicate computation.

4.3.2 Microscopic loops

The following theorem shows that loops of $\mathbb{P}^n_{LC}$ do not shrink to points as $n \to \infty$.

Theorem 8. *Any subsequential limit of $\mathbb{P}^n_{LC}$ as $n \to \infty$ is supported on Ω, that is*

$$\limsup_n \mathbb{P}^n_{LC}\,(\textit{there is a loop of area} \leq \alpha) = o(\alpha)\,. \tag{4}$$

Proof. Let us argue by contradiction. If we suppose that (4) is false, it means that there exists $p > 0$ and arbitrarily small values of $\delta > 0$ such that for n large enough (depending on δ), we have

$$\mathbb{P}^n_{LC}\left(\text{there is a loop with area} \leq \delta^2\right) > p\,.$$

Hence, using the Markov property and the cycle-popping algorithm, we show that there exists $0 < q < p$ and arbitrarily small values of $\delta > 0$ such that for n large enough (depending on δ), we have

$$\mathbb{P}^n_{LC}\left(\text{there are } 1/\delta \text{ loops with area} \leq \delta^2\right) > q^{1/\delta}\,.$$

However, by Lemma 27 below, there exists $0 < \xi < q$ such that for any k and n large enough, we have

$$\mathbb{P}^n_{LC}\,(\text{there are} \geq k \text{ loops}) \leq \xi^k\,.$$

By taking $k = 1/\delta$ we obtain that for arbitrarily small values of $\delta > 0$, there is a large enough n such that

$$0 < q^{1/\delta} < \mathbb{P}^n_{LC}\left(\text{there are } 1/\delta \text{ loops with area} \leq \delta^2\right) \leq \xi^{1/\delta}\,.$$

This yields a contradiction since the right-hand side of the last equation tends to zero faster than the left-hand side when $\delta \to 0$. $\square$

4.3.3 Resampling and tightness

We will show tightness of the sequence of measures $\mathbb{P}^n_{inc}, \mathbb{P}^n_{LC^0}$ and $\mathbb{P}^n_{LC}$. This will yield the existence of subsequential limits.

We show in fact that the scaling limits of the macroscopic loops are absolutely continuous with respect to SLE_2. Let γ be a loop in a $\mathbb{P}^n_{LC^0}$-,$\mathbb{P}^n_{LC}$- or $\mathbb{P}^n_{inc}$-random CRST, and $a, b \in \gamma$ distinct points on it. Let $\gamma[a, b]$ be the part of γ counterclockwise between a and b, and $\gamma[b, a]$ the complementary part. If we erase $\gamma[a, b]$, we can define a new loop γ' by taking a LERW from a to b in the domain defined by $\Sigma \cup \gamma[b, a]$, with wired boundary conditions on $\gamma[b, a]$, and conditioning on the LERW to end at b. The union of this LERW and $\gamma[b, a]$ is the simple closed curve γ'. By the sampling algorithm, this curve γ' is absolutely continuous with respect to $\gamma[a, b]$, with probability ratio given by the ratio of the cycle weights. If a and b are close to each other, the LERW from a to b will with high probability not exit a small ball around $[a, b]$. Hence the cycle weight of γ' will be close to that of γ. Wilson's algorithm thus shows that this is a fair sample of $\mathbb{P}$ conditioned on $\gamma[b, a]$. Thus the LERW from a to b with the appropriate boundary conditions is absolutely continuous with respect to $\gamma[a, b]$, with a bound independent of mesh size $1/n$.

This proves that in the scaling limit, the loops are locally absolutely continuous with respect to the scaling limit of LERW, which was shown in [10] to be SLE_2.

In particular the scaling limit is supported on simple curves (recall that SLE_2 curves are simple, and note that this is a local property).

We also need to check that for $\mathbb{P}^n_{inc}$ and $\mathbb{P}^n_{LC}$ the number of loops doesn't blow up as $n \to \infty$. This is a consequence of Lemma 27 below.

Theorem 9. *The sequence $(\mathbb{P}^n_{inc})_{n\geq 0}$ is tight on Ω.*

The above argument also shows that the sequence $(\mathbb{P}^n_{LC^0})_{n\geq 0}$ is tight, provided we allow for the possibility that the curve shrinks to a point. We need to add to Ω_1 a copy of Σ whose points represent the constant maps of S^1 to that point. The metric naturally extends to this augmented space Ω_1^*, so that the limit of a sequence of curves shrinking to a point is the constant curve at that point. Let Ω_1^* be this augmented space.

Theorem 10. *The sequence $\left(\mathbb{P}^n_{LC^0}\right)_{n\geq 0}$ is tight on Ω_1^*.*

It may be that $\left(\mathbb{P}^n_{LC^0}\right)_{n\geq 0}$ is tight on Ω_1 but we have not proved that its curves do not shrink to points as $n \to \infty$.

By Theorem 8 above, $(\mathbb{P}^n_{LC})_{n\geq 0}$ is tight on Ω, that is, its curves do not shrink to points.

Theorem 11. *The sequence $(\mathbb{P}^n_{LC})_{n\geq 0}$ is tight on Ω.*

4.4 Probabilities of cylindrical events

4.4.1 Flat connections

Let $x_1, \ldots, x_k \in \Sigma$. Let $\mathcal{M}$ be the space of flat SU_2-connections modulo gauge transformations on $\Sigma \setminus \{x_1, \cdots, x_k\}$. Such a flat connection is determined uniquely by a homomorphism from π_1 of the surface into SU_2.

The fundamental group $\pi_1(\Sigma \setminus \{x_1, \ldots, x_k\})$ is a free group F_m on $m = 2g+k-1+b$ letters, where g is the genus and b the number of boundary components of Σ. Thus a flat connection is determined by m arbitrary elements of SU_2, one for each generator of π_1. Let ν be the canonical measure on $\mathcal{M}$: the image of Haar measure on SU_2^m.

4.4.2 Trivalent graph

A useful device, see [4], is to define a trivalent graph H in Σ (unrelated to $\mathcal{G}$) with a single puncture of Σ in each face, so that H is a deformation retract of Σ. We can thus think of Σ as a ribbon graph structure on H.

Recall that a *finite lamination* L of Σ is an isotopy class of a finite number of simple curves. A lamination L retracts to a multicurve (with "multiplicity") on H; L is determined by, for each edge of H, a nonnegative integer giving the number of strands of a minimal representative of L retracted onto that edge. These integers satisfy the conditions that at each vertex of H the sum of the three integers is even and the three integers satisfy the triangle inequality: see Figure 8. Moreover any set of integers satisfying these two conditions arises from a unique lamination.

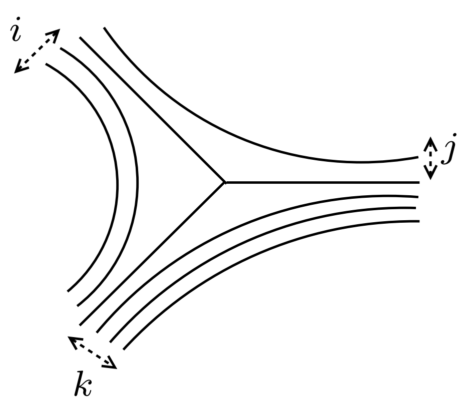

Figure 8: Strands of the lamination at a vertex of the trivalent graph; the integers i, j, k satisfy $i + j + k = 0 \mod 2$ and $|i - j| \leq k \leq i + j$.

We define a partial order on laminations: $L \leq L'$ if the integers on edges of H associated to L are all less than or equal to those of L'. We define the *complexity* $n(L)$ of L to be the sum of these integers.

An SU_2-connection on H determines a flat connection on Σ. After gauge transformation one can take the SU_2-connection to be the identity on all edges of a spanning tree of H.

4.4.3 Integrals over the space of flat connections

Given a finite weighted graph $\mathcal{G}$ embedded in Σ, and $\Phi \in \mathcal{M}$ define

$$Z(\Phi) = \sum_{\text{CRSFs}} \prod_{\text{edges}} c(e) \prod_{\text{cycles } \gamma} \left(2 - \mathrm{Tr}(\omega_\gamma)\right), \tag{5}$$

where ω_γ is the monodromy of the connection Φ around the cycle γ, for either choice of its orientation (and any starting vertex). The function Z is a real-valued function on $\mathcal{M}$ the space of flat connections modulo gauge.

We denote

$$Z_0 = \sum_{\text{inc. CRSFs}} \prod_{\text{edges}} c(e),$$

the partition function for all incompressible CRSFs, without cycle weight. (This is *not* the same as $Z(\mathrm{Id})$, which is zero.)

For a flat connection Φ we may rewrite (5) as

$$Z(\Phi) = \sum_L X_L T_L(\Phi),$$

where the sum is over finite laminations L, where X_L is the conductance-weighted sum of CRSFs whose cycles are isotopic to L, and

$$T_L(\Phi) = \prod_{\gamma \in L} (2 - \mathrm{Tr}(\omega_\gamma))$$

is a real-valued function on $\mathcal{M}$.

Fock and Goncharov [4] proved that, seen as real-valued functions on $\mathcal{M}$, the functions $S_L = \prod_{\gamma \in L} \mathrm{Tr}(\omega_\gamma)$ are linearly independent and generate the vector space of polynomial functions on $\mathcal{M}$, when L runs through all finite laminations. Hence the functions $T_L = \prod_{\gamma \in L} (2 - \mathrm{Tr}(\omega_\gamma))$ are also linearly independent, generate the same vector space and, when the bases are ordered by increasing number of cycles, the change of basis matrix $M_{S,T}$ is an invertible infinite triangular matrix.

Choose an ordering of the T_L consistent with the partial order on laminations defined above. Let $\mathbf{P} = \{P_L(\Phi)\}$ be the Gram-Schmidt orthonormalization of the $\mathbf{T} = \{T_L(\Phi)\}$ with respect to this ordering and with respect to the inner product $\langle f, g\rangle = \int_{\mathcal{M}} fg\, d\nu$. Let $\mathbf{A} = (A_{L,L'})$ be the infinite lower-triangular matrix such that $\mathbf{P} = \mathbf{A}\mathbf{T}$.

Recall the linear operator Δ_Φ on the total space $\mathbb{C}^{2|V|}$ of the $\mathbb{C}^2$-bundle on $\mathcal{G}$.

Theorem 12 ([8])**.** *We have*

$$Z(\Phi)^2 = \det(\Delta_\Phi).$$

One can extract the coefficients of any desired lamination L as follows.

Lemma 13. *For any cylindrical event E_L, we have*

$$\mu_{inc}(E_L) = \sum_{L' \geq L} A_{L',L} \int_{\mathcal{M}} \frac{Z(\Phi)}{Z_0} P_{L'}(\Phi) d\nu.$$

This sum is finite for any finite graph.

Proof. The probability of E_L is X_L/Z_0. Write

$$Z(\Phi) = \sum_L X_L T_L = \mathbf{X}^t \mathbf{T}(\Phi) = \mathbf{X}^t \mathbf{A}^{-1} \mathbf{P}(\Phi)\,.$$

Since $\mathbf{P}$ is orthonormal, we have $\int_{\mathcal{M}} \mathbf{P}^t \mathbf{P} d\nu = \mathrm{Id}$. Hence,

$$\begin{aligned}
\int_{\mathcal{M}} Z\mathbf{P}^t\mathbf{A} d\nu &= \int_{\mathcal{M}} \mathbf{X}^t \mathbf{A}^{-1} \mathbf{P}(\Phi)\mathbf{P}^t(\Phi)\mathbf{A} d\nu \\
&= \mathbf{X}^t \mathbf{A}^{-1} \left(\int_{\mathcal{M}} \mathbf{P}(\Phi)\mathbf{P}^t(\Phi)\, d\nu \right) \mathbf{A} = \mathbf{X}^t\,.
\end{aligned}$$

Hence,

$$\mathbf{X} = \int_{\mathcal{M}} Z\mathbf{A}^t \mathbf{P} d\nu\,.$$

Each entry X_L is the integral

$$X_L = \sum_{L' \geq L} A_{L',L} \int_{\mathcal{M}} Z(\Phi) P_{L'}(\Phi) d\nu.$$

Dividing by Z_0 we obtain the result. $\square$

4.5 Convergence in the flat case

We first consider Φ to be a flat connection. Associated to this is a measure μ^n_Φ on incompressible CRSFs. Since μ^n_Φ has a density (independent of the graph) with respect to μ^n_{inc}, it suffices to show that this latter converges.

The main tool is the following convergence result. Let $x_1, \dots, x_k$ be points of Σ and B_j a small ball around x_j. Let Φ be a flat connection on $\Sigma \setminus \{B_1 \cup \dots \cup B_k\}$.

Theorem 14 ([9])**.** *There exists a function $F \in \mathrm{L}^2(\mathcal{M})$ depending only on the conformal type of the surface $\Sigma \setminus \{B_1 \cup \dots \cup B_k\}$ such that*

$$\frac{Z(\Phi)}{Z_0} \to F(\Phi)\,.$$

The proof follows from Theorem 12 and the convergence of discrete harmonic functions to continuous ones.

Theorem 15. *Let $\mathcal{G}_n$ be a sequence of graphs with mesh size going to zero and conformally approximating a non-simply connected surface Σ. The measures $\mathbb{P}^n = \mathbb{P}^n_{inc}$ which are uniform on incompressible CRSFs of $\mathcal{G}_n$ converge as $n \to \infty$, and the limit is a conformally invariant measure $\mathbb{P}_{inc}$ supported on sets of pairwise disjoint, noncontractible simple loops on Σ.*

In [8] the homotopy classes on Σ of the noncontractible loops were shown to have a conformally invariant limit distribution.

Proof. Take points $z_1, \dots, z_k$ in Σ, take $\delta > 0$ small, and for each i let B_i be the ball of radius δ around z_i. Let $\mathcal{G}^n_B = \mathcal{G}^n \setminus \{B_1 \cup \dots \cup B_k\}$ and $\mathbb{P}^n_B$ the associated measure on multiloops of incompressible CRSFs on $\mathcal{G}_n$ whose loops stay in $\mathcal{G}^n_B$. A loop on $\Sigma \setminus \{B_1 \cup \dots \cup B_k\}$ is said to be *peripheral* if it is isotopic to one of the boundary curves ∂B_i. Laminations on $\mathcal{G}^n_B$ may contain peripheral loops, but we are interested in those without peripheral loops.

Up to errors uniform in n and tending to zero with δ,

$$\mathbb{P}^n(E_L) = \mathbb{P}^n_B\left(E_L \mid \text{no peripheral loops}\right).$$

This follows from the sampling algorithm since removing one or more very small holes does not change the distribution of the LERW.

It remains to show that $\lim_{n\to\infty} \mathbb{P}^n_B(E_L)$ exists and depends only on the conformal type of the domain $\Sigma \setminus \{B_1 \cup \dots \cup B_k\}$; the conditioning multiplies by a non-zero constant which is also a conformal invariant quantity and independent of L.

Consider L a finite lamination in $\Sigma \setminus \{B_1 \cup \cdots \cup B_k\}$, which has no peripheral curves. By Lemma 13 the probability $P_B^n(E_L)$ is given by a sum over $L' \geq L$ of integrals over $\mathcal{M}$ of $Z(\Phi)/Z_0$ times a function $P_{L'}(\Phi)$ independent of n.

By Theorem 14 the integrand $Z(\Phi)/Z_0$ converges. Moreover, it is bounded independently of n as follows. We write

$$Z_\Phi/Z_0 = \sum_L X_L/Z_0 T_L \,. \tag{6}$$

First $T_L \leq e^{O(n(L))}$ because $2 - \mathrm{Tr}(\omega)$ is uniformly bounded by a constant over $\mathcal{M}$. Secondly, by Lemma 27 below, $X_K \leq e^{-cn(K)} Z_0$ for arbitrarily small $c > 0$. Since there are at most $n(K)^M$ laminations L with complexity $n(L) = n(K)$, where M is the number of edges of H, the sum (6) is bounded by a converging series. Hence by bounded convergence, for each L' the integral $\int_{\mathcal{M}} \frac{Z(\Phi)}{Z_0} P_{L'} d\nu$ converges.

We need to show that the sum over L' converges. We write

$$\langle \sum_K X_K T_K, P_{L'} \rangle = \langle \sum_{K \geq L'} X_K T_K, P_{L'} \rangle \tag{7}$$

since $\langle T_K, P_{L'} \rangle = 0$ unless $K \geq L'$. We also have $|\langle T_K, P_{L'} \rangle| \leq e^{O(n(K))}$ by the Cauchy-Schwarz inequality. Using the above bounds the sum (7) is bounded by $e^{-c' n(L)} Z_0$ for some arbitrarily small $c' > 0$ and summing over L' now gives a convergent sum.

We now use a classical convergence argument [2]. We have shown that the probabilities of any cylindrical event converge. The cylindrical events form a family of sets which is stable under finite intersection. Since it generates the σ-field $\mathcal{B}$ it is a determining class, that is if two probability measures on $(\Omega, \mathcal{B})$ coincide on all the cylindrical events, then they are equal.

Since we furthermore have tightness by Theorem 9, the sequence of probability measures $\mathbb{P}^n$ admits subsequential limits by Prokhorov's theorem. More precisely, for any subsequence, there exists a subsubsequence which converges. Since its value on the cylindrical events is known, there is only one possible limit. Let us call it $\mathbb{P}_{inc}$. Now, since Ω is a metric space, this implies that the sequence μ_{inc}^n converges weakly to $\mathbb{P}_{inc}$. $\square$

As a corollary, we obtain the convergence of the measures μ_Φ^n for any flat connection Φ.

Corollary 16. *There exists a probability measure $\mathbb{P}_\Phi$ on $(\Omega, \mathcal{B})$ such that*

$$\mathbb{P}_\Phi^n \to \mathbb{P}_\Phi$$

in the sense of weak convergence.

Proof. It suffices to show the convergence of this sequence of measures on finite intersections of balls in Ω of small radius because this is a determining class for $\mathcal{B}$. For any curve γ, there is a small radius $r > 0$ such that its tubular r-neighborhood retracts onto γ. Any curve in this r-neighborhood and winding once around is isotopic to γ. On any such neighborhood the density $\prod_{\gamma \subset L} \left(2 - \omega_\gamma - \omega_\gamma^{-1}\right)$ is constant, hence the convergence follows by Theorem 15. $\square$

4.6 Curved case

The measures μ_{LC}, μ_{LC^0} converge in the following sense.

Theorem 17. *There exist probability measures $\mathbb{P}_{LC}, \mathbb{P}_{LC^0}$ on $(\Omega, \mathcal{B})$ and $(\Omega_1^*, \mathcal{B})$ respectively such that for any sequence $(\mathcal{G}_n)$ of graphs, geodesically embedded on Σ, with mesh size going to zero as $n \to \infty$ and conformally approximating Σ, the sequences of probability measures $\mathbb{P}^n_{LC}$ and $\mathbb{P}^n_{LC^0}$ converge weakly towards respectively $\mathbb{P}_{LC}$ and $\mathbb{P}_{LC^0}$.*

Proof. Let us approximate Σ by a polygonal surface Σ_ε, that is, with a surface which is flat except for conical singularities. A standard way to do this is to take a fine triangulation of the surface (with triangles of diameter at most ε and whose angles are bounded from below), and replace each triangle with the Euclidean triangle with the same edge lengths. As $\varepsilon \to 0$ the conformal structure on Σ_ε converges to that of Σ. (Indeed, there is a homeomorphism from Σ_ε to Σ which is $(1+o(1))$-biLipschitz.)

We furthermore suppose that the conical singularities are separated at distance $\delta(\varepsilon) \to 0$ with $\varepsilon = o(\delta)$.

Any graph embedded on Σ or Σ_ε can be embedded on Σ_ε or Σ with small distortion; furthermore a graph conformally approximating Σ_ε will have image on Σ conformally approximating Σ, and vice versa.

Let $z_1, \dots, z_k$ be the vertices of Σ_ε. The Levi-Civita connection on Σ_ε is a flat connection on $\Sigma_\varepsilon \setminus \{z_1, \dots, z_k\}$ and approximates the Levi-Civita connection on Σ, in the sense that the curvature enclosed by any loop is close for both connections (for small ε). Restricting to $\mathcal{G}_n$, this shows that $\mathbb{P}^n_{LC,\varepsilon}$ is close to $\mathbb{P}^n_{LC}$ (since cylinder events have close measure).

By Corollary 16, if we fix Σ_ε and take $\mathcal{G}_n$ embedded on Σ_ε and conformally approximating it as $n \to \infty$, the measures $\mathbb{P}^n_{LC,\varepsilon}$ converge as $n \to \infty$ to a limit $\mathbb{P}_{LC,\varepsilon}$. Similarly for $\mathbb{P}^n_{LC^0,\varepsilon}$.

For any $\varepsilon > 0$ the probabilities of the cylindrical events (away from the singularities) are determined by the function F of Lemma 14. These are expressed as integrals of the Green function on paths avoiding the singularities. The addition of a

singularity modifies the Green's function, and hence the probability, by a negligible amount. Since the singularities are at distance $\delta >> \varepsilon$ it follows that the probability measures $\mathbb{P}_{LC,\varepsilon}$ and $\mathbb{P}_{LC^0,\varepsilon}$ each form a Cauchy sequence.

By a diagonal argument, we can take $n \to \infty$ and $\varepsilon \to 0$ and conclude that $\mathbb{P}^n_{LC}$ and $\mathbb{P}^n_{LC^0}$ converge weakly as $n \to \infty$. □

From experimental simulations, the probability of getting two or more loops in μ_{LC} on the round sphere is on the order of one percent. Hence measures μ_{LC} and μ_{LC^0} are close for the total variation distance.

The measure $\mathbb{P}_{LC^0}$ is a limit of $\mathbb{P}_{LC}$ when the metric is scaled by a factor $t \to 0$. Hence it is not obvious that Theorem 8 implies that $\mathbb{P}_{LC^0}$ also is supported on macroscopic loops. However, we conjecture it to be true.

Conjecture 18. $\mathbb{P}_{LC^0}$ *is supported on* Ω_1 *that is*

$$\mathbb{P}_{LC^0}\left(\textit{the area of the loop is zero}\right) = 0\,.$$

4.7 A comment and two applications

The convergence result of probability measures $\mu_{c,\alpha}$ on CRSFs is actually more general and can be adapted for a wide range of functions α on the cycles, not necessarily coming from connections. This is due to the fact that the crucial convergence argument is made for the uniform measure on incompressible CRSFs (Theorem 15). We now state two applications of the previous theorems.

As an application of Theorem 17, we obtain a positive answer to Conjecture 1 of [7]. Let $\mathcal{G}_n$ be a rectilinear approximation of a planar Jordan domain D (see [7] for a definition). Let A denote the combinatorial area (number of faces) of the cycle of a cycle-rooted spanning tree of $\mathcal{G}_n$. Let $\theta = A/n^2$ be the Euclidean area of the cycle. We denote by $\mathbb{E}^n_{\text{unif}}$ the expectation with respect to the uniform measure on CRSTs on $\mathcal{G}_n$.

Corollary 19. *For* $k \geq 2$ *there exists* $a_k(D) > 0$ *such that*

$$\mathbb{E}^n_{unif}\left(A^k\right) = a_k(D) n^{2k-2}(1 + o(1))\,.$$

Proof. Let $k \geq 2$. We have

$$\begin{aligned}
\mathbb{E}^n_{\text{unif}}\left(A^k\right) &= \mathbb{E}^n_{LC^0}\left(A^{k-2}\right)\mathbb{E}^n_{\text{unif}}\left(A^2\right) \\
&= n^{2k-4}\mathbb{E}^n_{LC^0}\left(\theta^{k-2}\right)\mathbb{E}^n_{\text{unif}}\left(A^2\right) \\
&= n^{2k-2}C(D)|D|\mathbb{E}_{LC^0}\left(\theta^{k-2}\right)(1 + o(1))\,,
\end{aligned}$$

where the last equality follows from Theorem 8 of [7] (here $C(D)$ is a constant) and the weak convergence of $\mathbb{P}^n_{LC^0}$ to $\mathbb{P}_{LC^0}$. Since θ is bounded and for $\mathbb{P}_{LC^0}$ is with positive probability non-zero by Theorem 7, the limit $\mathbb{E}_{LC^0}\left(\theta^{k-2}\right)$ is a positive real. The corollary is proved by taking $a_k(D) = C(D)|D|\mathbb{E}_{LC^0}\left(\theta^{k-2}\right)$. □

As an application of Theorem 15, let us state the following corollary.

Corollary 20. *Consider a uniform spanning forest on an annulus-graph wired on its boundary. Then the simple closed curve separating the two connected components has a conformally invariant limit which is given by the measure $\mathbb{P}_{inc}$ conditional on having only one component.*

Proof. This follows from the fact that the dual of a wired essential forest on the annulus is a uniform incompressible CRST on the annulus with free boundary conditions. The measure is thus given by μ^n_{inc} conditional on having one loop. The convergence follows from Theorem 15. □

Figure 9 shows a sample of this interface between the two tree components of a wired uniform spanning forest on the annulus.

5 Properties of the measures

In this section, we mention a few properties of the measures on CRSFs on surfaces we have been considering. This could help give a characterization of these measures.

5.1 Markov property

CRSFs on surfaces satisfy the following *spatial Markov property.* Consider a graph $\mathcal{G}$ embedded in a compact oriented surface Σ of positive genus. Let $\alpha : \Omega_1 \to \mathbb{R}_{>0}$ be any positive weight function on the cycles of $\mathcal{G}$.

Let $\{\gamma\} = \{\gamma_1, \dots, \gamma_k\}$ be a family of cycles in $\mathcal{G}$ which separate Σ in a number of connected components $\Sigma_1, \dots, \Sigma_r$, with $r \leq k$. For $i = 1, \dots, r$, denote by $\mathcal{G}_i$ the intersection of $\mathcal{G}$ and the closure of Σ_i (*i.e.* Σ_i along with its boundary) and by $\partial\mathcal{G}_i$ the boundary cycles.

An *essential CRSF with Dirichlet boundary conditions* on a graph with boundary is a spanning subgraph which is the union of disjoint unicycles not touching the boundary and a union of trees rooted on the boundary.

Let $\mathbf{\Gamma}$ be the random CRSF on $\mathcal{G}$ associated to the connection Φ. For each i, denote by $\mathbf{\Gamma}_i$ the random essential CRSF on the line bundle $\mathcal{G}_i$, with boundary $\partial\mathcal{G}_i$, associated to Φ_i obtained by restriction of Φ on $\mathcal{G}_i$.

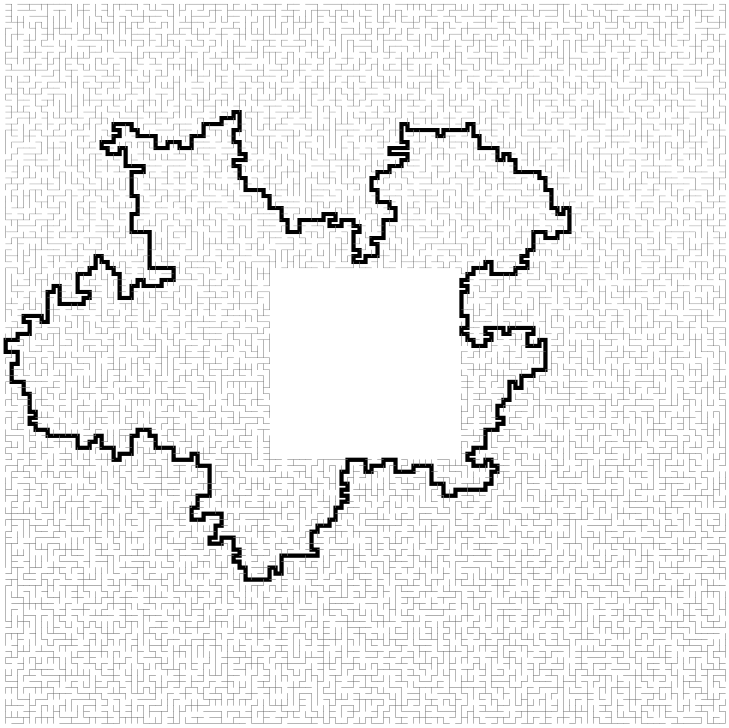

Figure 9: A uniform incompressible CRST on an annulus in the square grid

Proposition 21. *Conditional on $\{\gamma\} \subset \mathbf{\Gamma}$, the random CRSF $\mathbf{\Gamma}$ is equal in distribution to*

$$\{\gamma\} \sqcup_{i=1}^{r} \mathbf{\Gamma}_i \, .$$

Proof. The key argument is that, because of the assumption on $\{\gamma\}$, a CRSF Γ on $\mathcal{G}$ which contains $\{\gamma\}$ is the union $\{\gamma\} \sqcup_{i=1}^{r} \Gamma_i$ of $\{\gamma\}$ with essential CRSFs Γ_i on each one of the connected component Σ_i of $\Sigma \setminus \{\gamma\}$. The proposition then follows directly from the cycle popping algorithm. □

5.2 Restriction property

Let $D_1 \subset D$ be two planar Jordan domains. Let $\mathcal{G}$ be a finite graph approximation of D. Let Φ be a connection on a line bundle over $\mathcal{G}$. We denote $\det \Delta^D(\Phi)$ the line bundle Laplacian on $\mathcal{G}$ with connection Φ. For any subset S of the set of vertices of $\mathcal{G}$, we denote by $\det \Delta^D_S(\Phi)$ the line bundle Laplacian with Dirichlet boundary conditions on S, see [8].

Lemma 22. *For any finite set of simple non-intersecting curves $\{\gamma\} \subset D_1$, we have*

$$\frac{\mu_\Phi^{D_1}(\{\gamma\})}{\mu_\Phi^{D}(\{\gamma\})} = \left(\frac{\det \Delta^{D_1}(\Phi)}{\det \Delta^{D_1}_{\{\gamma\}}(\Phi)}\right) / \left(\frac{\det \Delta^{D}(\Phi)}{\det \Delta^{D}_{\{\gamma\}}(\Phi)}\right) .$$

5.3 Stochastic domination

Recall from Section 2 the notation $\mu_{c,\alpha}$ for measures that assign a CRSF Γ a probability proportional to $\prod_{e\in\Gamma} c(e) \prod_{\gamma\subset\Gamma} \alpha(\gamma)$.

Lemma 23. *Let $S_1 \subset S_2$ be two subgraphs of $\mathcal{G}$. Let $\mathbb{P}_1$ and $\mathbb{P}_2$ be essential CRSF measures $\mu_{c,\alpha}$ with Dirichlet boundary conditions on S_1 and S_2, respectively. For any curve $\gamma \in \mathcal{G} \setminus S_2$, we have*

$$\mathbb{P}_1(\gamma) \geq \mathbb{P}_2(\gamma) .$$

Proof. This follows from the definition of algorithm $\mathcal{A}$ described to sample these measures since it is less likely for LERW, started from a vertex outside S_2, to create loop γ before it hits S_2 than before it hits S_1. □

Note that in the case the measures come from a line bundle connection Φ, this implies

$$\frac{\det \Delta_{S_1\cup\gamma}(\Phi)}{\det \Delta_{S_1}(\Phi)} \geq \frac{\det \Delta_{S_2\cup\gamma}(\Phi)}{\det \Delta_{S_2}(\Phi)} ,$$

which is a non trivial potential theoretic consideration (which can be translated in terms of Dirichlet-to-Neumann map).

5.4 Sub-exponential tail for the distribution of the number of loops

Let $\mathbb{P}^n$ be one of the two sequences of measures $\mathbb{P}^n_{inc}$ or $\mathbb{P}^n_{LC}$. We show that the number of loops has tail decreasing faster than exponentially. This involves a computation from two specific examples, one for $\mathbb{P}_{LC}$ and one for $\mathbb{P}_{inc}$.

5.4.1 Wired boundary for $\mathbb{P}_{LC}$

Let $\mathcal{G}$ be a cylindrical square grid graph with height n and circumference n, obtained from a $(n+1)\times n$ square grid by identifying the left and right boundaries. We wire the upper and lower boundaries, that is, add an extra vertex connected to all points along the upper and lower boundary.

Let Φ be a line bundle connection on $\mathcal{G}$ with monodromy $q = e^{i\phi}$ around each square face, where $\phi = c/n^2$ for a constant c. Thus the total curvature is c.

We can realize Φ by putting parallel transport q^y from (x, y) to $(x+1, y)$ (indices cyclic) and the identity on vertical edges. We compute the partition function $Z(\Phi)$ of μ_Φ, the measure on essential CRSFs of $\mathcal{G}_n$ (CRSFs with wired boundary).

Theorem 24. $\frac{Z_\Phi}{Z_I} = 1 + O(c^2)$.

Corollary 25. *With probability* $1 - O(c^2)$ *there are no loops.*

Proof of Theorem. The matrix of the Dirichlet Laplacian Δ_Φ commutes with horizontal translation (rotation of the cylinder). The action of Δ_Φ on the $e^{i\theta}$-eigenspace (here $\theta = \frac{2\pi k}{n}$ for $k = 0, \ldots, n-1$) for horizontal translation is

$$\Delta_\Phi(\theta) = \begin{pmatrix} 4 - 2\cos\theta & -1 & & \\ -1 & 4 - 2\cos(\theta + \phi) & -1 & \\ & -1 & \ddots & -1 \\ & & -1 & 4 - 2\cos(\theta + (n-1)\phi) \end{pmatrix}.$$

For small ϕ we have $\Delta_\Phi = \Delta_I + S$ for a small diagonal perturbation S, so

$$\det \Delta_\Phi = \det \Delta_I \det(I + S\Delta_I^{-1}).$$

Keeping terms of order up to ϕ^2 gives

$$\frac{Z_\Phi}{Z_I} = \frac{\det \Delta_\Phi}{\det \Delta_I} =$$

$$= \prod_\theta \left(1 + \sum_{j=1}^n S_{jj}(\theta)\Delta_{I,j,j}^{-1}(\theta) + \sum_{i<j} S_{i,i}(\theta)S_{j,j}(\theta)\left[\Delta_{I,i,i}^{-1}(\theta)\Delta_{I,j,j}^{-1}(\theta) - \Delta_{I,i,j}^{-1}(\theta)\Delta_{I,j,i}^{-1}(\theta)\right]\right).$$

Here the product is over $\theta = 2\pi k/n,\ k = 0, \ldots, n-1$, and

$$S_{jj}(\theta) = 2j\phi\sin\theta + j^2\phi^2\cos\theta + O(\phi^3).$$

The entries of $\Delta_I^{-1}(\theta)$ are written in terms of Chebyshev polynomials. It is convenient to define $\alpha > 0$ by $4 - 2\cos\theta = 2\cosh\alpha$. Then for $i \geq j$ we have

$$\Delta_{I,i,j}^{-1}(\theta) = \frac{\sinh(j\alpha)\sinh((n+1-i)\alpha)}{\sinh\alpha\sinh(n+1)\alpha}.$$

(Since the matrix is symmetric this defines the case $i < j$ as well.)

If α is not small, $\Delta_{I,i,j}^{-1}(\theta) \approx e^{-|j-i|\alpha}/2\sinh\alpha$. For α small we can use instead the bound $\Delta_{I,i,j}^{-1}(\theta) \leq \frac{j(n+1-i)}{n+1}$.

Expanding out the product over θ we find that the first-order term in ϕ vanishes because $\sin\theta$ is odd. We are left with

$$\frac{Z_\Phi}{Z_I} = 1 + \phi^2(A_1 + A_2 + A_3) + O(\phi^3),$$

where

$$A_1 = \sum_\theta \sum_{j=1}^{n} j^2 \frac{\sinh(j\alpha)\sinh((n+1-j)\alpha)}{\sinh\alpha\sinh(n+1)\alpha}\cos\theta, \tag{8}$$

$$A_2 = \sum_\theta \sum_{j<k} 4jk\sin^2\theta \left[\Delta_{I,j,j}^{-1}(\theta)\Delta_{I,k,k}^{-1}(\theta) - \Delta_{I,j,k}^{-1}(\theta)\Delta_{I,k,j}^{-1}(\theta)\right],$$

and

$$\begin{aligned} A_3 &= \sum_{\theta_1<\theta_2} \sum_{j,k} 4jk\sin\theta_1\sin\theta_2\Delta_{I,j,j}^{-1}\Delta_{I,k,k}^{-1} \\ &= -\frac{1}{2}\sum_\theta \sum_{j,k} 4jk\sin^2\theta\Delta_{I,j,j}^{-1}\Delta_{I,k,k}^{-1}. \end{aligned}$$

Since $\phi = c/n^2$, it remains to show that $A_1 + A_2 + A_3 = O(n^4)$.

Summing A_2 and A_3 gives

$$A_2 + A_3 = -\sum_\theta \sum_{j<k} 4jk\sin^2\theta(\Delta_{I,j,k}^{-1})^2 - \frac{1}{2}\sum_\theta \sum_j 4j^2\sin^2\theta(\Delta_{I,j,j}^{-1})^2. \tag{9}$$

The second sum in (9) is $O(n^4)$. The first is $O(n^4)$ if we restrict to $|\theta| > \varepsilon$ for a constant ε. Similarly the sum (8) is $O(n^4)$ under this restriction. It now remains to consider the case of $|\theta| < \varepsilon$.

Consider first A_1. The sum over j can be computed explicitly as a power series; the leading term is of order n^3 and is

$$\frac{n^3}{6}\frac{\cos\theta\cosh(n+1)\alpha}{\sinh\alpha\sinh(n+1)\alpha}.$$

When θ is not $O(1/n)$, we have $\frac{\cosh(n+1)\alpha}{\sinh(n+1)\alpha} \approx 1$, and $\sinh\alpha = \theta + O(\theta)^2$ so the leading term is

$$n^3\left(\frac{1}{6\theta} + O(1)\right).$$

When $\theta = 2\pi j/n$ and j is of order 1, $\frac{\cosh(n+1)\alpha}{\sinh(n+1)\alpha}$ is of constant order and so the terms for small θ contribute a negligible amount to the final sum.

Similarly the sum over j in (9) can be computed explicitly (we resorted to using a computer algebra program); the leading term is of order n^3 and is

$$-\frac{n^3}{3}\frac{\sin^2\theta}{(e^{2\alpha}-1)\sinh\alpha} = n^3\left(-\frac{1}{6\theta} + O(1)\right)$$

for θ not on the order of $1/n$.

The leading terms in A_1 and $A_2 + A_3$ thus cancel and the sum is $n^3\sum_\theta O(1) = O(n^4)$.

Thus

$$\frac{Z_\Phi}{Z_I} = 1 + \frac{c^2}{n^4}O(n^4) = 1 + O(c^2).$$

$\square$

A similar computation holds for a cylinder with aspect ratio $m/n \to \tau$. By universality, the results holds for any graph conformally approximating the curved annulus. If the curvature is not constant the result still holds since the measures for the constant curvature and variable curvature case are absolutely continuous with bounded Radon-Nikodym derivative between them.

A rectangle with wired boundary can be obtained from an annulus with wired boundary by wiring the vertices along a zipper connecting the two boundary components. Since wiring increases the probability of having no loops, we can conclude that for a rectangle with total curvature c, the probability of having no loops is also $1 - O(c^2)$. By universality and conformal invariance this holds for any disk and similarly for any planar Riemann surface.

5.4.2 Wired boundary for $\mathbb{P}_{inc}$

As above let $\mathcal{G}$ be a cylindrical square grid graph with height m and circumference n, obtained from a $(n+1)\times m$ square grid by identifying the left and right boundaries.

Let Φ be a flat line bundle on $\mathcal{G}$ with monodromy z around the boundary. We can realize Φ by putting parallel transport u from (x, y) to $(x+1, y)$ (indices cyclic) and the identity on vertical edges, where $u^n = z$.

We compute the partition function $Z(z)$ of μ_Φ, the measure on incompressible essential CRSFs of $\mathcal{G}$ (CRSFs with wired boundary).

Theorem 26. *As $n, m \to \infty$ with $n/m \to \tau$, the probability generating function for the number of loops tends to*

$$p_\tau(X) = \prod_{j=1}^{\infty} \frac{q^j + q^{-j} - 2 + X}{q^j + q^{-j} - 1},$$

where $q = e^{\pi\tau}$.

In particular the probability that there are no loops tends to a nonzero constant. The analogous computation with free boundary was done in [8].

Proof. By [8] we have

$$Z(z) = \prod_{\lambda} \mathrm{Ch}_n(\lambda + 2) - z - 1/z$$

where Ch_n is a variant of the Chebyshev polynomial defined by $\mathrm{Ch}_n(\alpha + 1/\alpha) = \alpha^n + \alpha^{-n}$ and λ runs over the eigenvalues of the line graph Laplacian with wired endpoints, that is, eigenvalues of the $m \times m$ matrix

$$\begin{pmatrix} 2 & -1 & \\ -1 & \ddots & -1 \\ & -1 & 2 \end{pmatrix}.$$

These eigenvalues are $2 - 2\cos\frac{\pi j}{m+1}$ for $j = 1, \ldots, m$.

Setting $X = 2 - z - 1/z$ in $Z(z)$, the power of X records the number of loops in the configuration. We have

$$Z(z) = \prod_{\lambda} \mathrm{Ch}_n(\lambda + 2) - 2 + X,$$

and the probability generating function for the number of loops is

$$\prod_{j=1}^{m} \frac{\mathrm{Ch}_n(4 - 2\cos\frac{\pi j}{m+1}) - 2 + X}{\mathrm{Ch}_n(4 - 2\cos\frac{\pi j}{m+1}) - 1}. \tag{10}$$

When n is large, $\mathrm{Ch}_n(w)$ is exponentially large in n for real w unless w is near 2. Thus the X dependence of (10) arises from j near 0. Letting $n, m \to \infty$ with $n/m \to \tau$ this expression is asymptotic to (with $q = e^{\pi\tau}$)

$$\prod_{j=1}^{\infty} \frac{q^j + q^{-j} - 2 + X}{q^j + q^{-j} - 1}.$$

□

5.4.3 Loop number lemma

Lemma 27. *For any $0 < \xi < 1$, there exists N and K such that for all $n \geq N$ and for all $k \geq K$, we have*

$$\mathbb{P}^n\left(\textit{there are at least } k \textit{ loops}\right) \leq \xi^k .$$

Proof. Each loop created during the performance of the cycle-popping algorithm either disconnects the surface or decreases the rank of the first homology. By the Markov property, the law of the conditional CRSF is obtained by independently sampling in each of the connected components.

For $\mathbb{P}^n_{LC}$, suppose that we have created k closed curves. Each of the complementary components is either planar or nonplanar. Each new loop found has a positive probability of being macroscopic, that is, either reduces the rank of the first homology or removes definite area from both resulting components, by Theorem 7. Thus for k sufficiently large, the curvature enclosed c on either side will be small, hence with definite probability $1-p = 1-O(c^2)$ the two sides will contain no further loops (by Corollary 25 above). By taking k large enough, and mesh size $1/n$ small enough, the curvature enclosed c can be taken as small as needed such that $p < \xi$.

Thus the probability of an extra loop eventually decays exponentially with rate less than ξ.

A similar argument works for $\mathbb{P}_{inc}$ using Theorem 26 instead of Corollary 25. □

In fact we conjecture that the decay of the probability of the number of loops decays faster than exponentially since each new loop subdivides a region into regions each having curvature smaller by a definite fraction. This would imply a decay of order $\xi^k/k!$. We conjecture an even stronger decay as follows.

Conjecture 28. *Let $\mathbb{P}$ be one of the measures $\mathbb{P}_{LC}$ or $\mathbb{P}_{inc}$. There exists $0 < \xi < 1$ such that for all $k \geq 1$, we have*

$$\mathbb{P}\left(\textit{there are at least } k \textit{ loops}\right) \leq \xi^{k^2} .$$

6 Questions

1. Can the measures $\mathbb{P}_{inc}$, $\mathbb{P}_{LC^0}$, and $\mathbb{P}_{LC}$ be defined directly instead of via limits of CRSF measures? For example via a stochastic differential equation, like a variant of SLE_2 defined on Riemann surfaces subject to some potential depending on the metric?

2. On the round sphere for the measure μ_{LC^0}, can one use the connection Laplacian to say more about the shape of the cycle, as is done in [7] in the flat case? By Corollary 19 above, the expected area of a $\mathbb{P}_{LC^0}$-random loop is $\frac{a_3(D)}{C(D)|D|}$, where $a_3(D)n^4$ is the leading order of the third moment of the combinatorial area of the loop of a uniform cycle-rooted spanning tree on $\mathcal{G}_n$. Can one compute $a_3(D)$? Is there a number theoretic interpretation?

3. What can be said about the Gaussian Free Field associated to the line bundle Laplacian? Are our loop models related to this GFF? In particular, for a choice of an infinite curvature, we expect to obtain loops at all scales.

4. What is the right framework for the study of $\mathbb{P}_{LC}$? What is the distribution of the number of loops?

5. Is there a probabilistic interpretation for the coefficients of the triangular matrix $\mathbf{A}$ defined in Section 4.4.3?

References

[1] D. Aldous, *On the time taken by random walks on finite groups to visit every state*, Wahrsch. Verw. Gebeite 62 (1983), 361–374.

[2] P. Billingsley, *Convergence of probability measures*, second edition, Wiley Series in Probability and Statistics: Probability and Statistics, Wiley & Sons, New York, 1999.

[3] R. Burton, R. Pemantle, *Local characteristics, entropy, and limit theorems for spanning trees and domino tilings via transfer impedances*, Ann. Prob. 21 (3) (1993), 1329–1371.

[4] V. Fock, A. Goncharov, *Moduli spaces of local systems and higher Teichmüller theory*, Publ. Math. Inst. Hautes Études Sci. No. 103 (2006), 1–211.

[5] R. Forman, *Determinants of Laplacians on graphs*, Topology 32 (1) (1993), 35–46.

[6] J. B. Hough, M. Krishnapur, Y. Peres, B. Virág, *Determinantal processes and independence*, Prob. Surveys (3) (2006), 206–229.

[7] A. Kassel, R. Kenyon, W. Wu, *On the uniform cycle-rooted spanning tree in $\mathbb{Z}^2$*, preprint, arXiv:1203.4858.

[8] R. W. Kenyon, *Spanning forests and the vector bundle Laplacian*, Ann. Prob. 39 (5) (2011), 1983–2017.

[9] R. W. Kenyon, *Conformal invariance of loops in the double dimer model*, preprint, arXiv:1105.4158.

[10] G. Lawler, O. Schramm, W. Werner, *Conformal invariance of planar loop-erased random walks and uniform spanning trees*, Ann. Probab. 32 (2004), no. 1B, 939–995.

[11] D. B. Wilson, *Generating random spanning trees more quickly than the cover time*, Proceedings of the Twenty-eigth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996), 296-303, ACM, New York, 1996.